A Constrained Optimization Algorithm for Total Energy Minimization in Electronic Structure Calculation *

Chao Yang [*,1] Juan C. Meza [1] Lin-Wang Wang [1]

Abstract

A new direct constrained optimization algorithm for minimizing the Kohn-Sham (KS) total energy functional is presented in this paper. The key ingredients of this algorithm involve projecting the total energy functional into a sequences of subspaces of small dimensions and seeking the minimizer of total energy functional within each subspace. The minimizer of a subspace energy functional not only provides a search direction along which the KS total energy functional decreases but also gives an optimal "step-length" to move along this search direction. A numerical example is provided to demonstrate that this new direct constrained optimization algorithm can be more efficient than the self-consistent field (SCF) iteration.

Key words: electronic structure calculation, total energy minimization, nonlinear eigenvalue problems, constrained optimization

1 Introduction

One of the fundamental problems in electronic structure calculations is to minimize the Kohn-Sham (KS) total energy functional with respect to electron wave functions. Currently, the most widely used approach for solving this

[*] Corresponding author
  Email address: `cyang@lbl.gov` (Chao Yang).
[1] Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

minimization problem is to apply the so called *Self Consistent Field* (SCF) iteration to the nonlinear equation derived from the first order necessary optimality condition. In each SCF iteration, one must compute approximations to a few smallest eigenvalues and their corresponding eigenvectors of a large matrix (Hamiltonian).

Methods for minimizing the total energy directly have been examined in the past (Kresse & Furthmüller 1996a; Gillan 1989; Teter 1989; Payne et al. 1992; Pfrommer et al. 1999; Voorhis & Head-Gordon 2002; Bendt & Zunger 1982). These methods construct a search direction based on the gradient of the total energy, and perform some type of line search along that direction in order to determine an optimal step length. The difficulty with these approaches is that the line search strategy must take into account the orthonormality constraints imposed on the wave functions. Due to this difficulty, the minimization algorithms developed in the past could be slower than SCF by a factor of 1.5 to 10 (Kresse & Furthmüller 1996a).

In this paper, we show that direct minimization of the KS functional can be made more effective than the SCF algorithm. Instead of using a simple line search strategy, we project the total energy into a subspace from which an optimal search direction and step length are determined simultaneously by solving a smaller nonlinear eigenvalue problem. In our approach, the optimal wave functions are determined simultaneously rather than "band-by-band". The orthonormality constraint of these wave functions is automatically satisfied by the solution to the small nonlinear eigenvalue problem.

2 Mathematical background

In this section, we establish the mathematical notation required to describe a new constrained optimization algorithm. We begin with the continuous formulation of the optimization problem, and establish the finite-dimensional analog in terms of linear algebra notation.

Following the notation established in (Goedecker & Scuseria 2003), we denote the i-th electron wave function by $\psi_i(r)$, where r represents the (three-dimensional) spatial coordinates. We denote the electron charge density by $\rho(r)$. It is defined by

$$\rho(r) = \sum_{i=1}^{k} |\psi_i(r)|^2,$$

where k is the number of occupied states.

The KS total energy functional consists of several components, i.e.,

$$E_{total}(\{\psi_i\}) = E_{kinetic} + E_{ion} + E_H + E_{XC}, \tag{1}$$

where $E_{kinetic}$ is the kinetic energy, E_{ion}, E_H and E_{XC} are potential energies induced by the electron-ion interaction (ionic potential), the electron-electron interaction (Hartree potential) and the exchange correlation potential respectively.

The kinetic energy of the atomistic system is defined by

$$E_{kinetic} = -\frac{1}{2} \sum_{i=1}^{k} \int \bar{\psi}_i(r) \nabla^2 \psi_i(r) dr,$$

where ∇^2 is the Laplacian operator, and $\bar{\psi}$ is the complex conjugate of ψ.

The ionic potential energy consists of a local and a non-local term. The local term can be expressed by

$$E_{ion(local)} = \int V_{ion}(r)\rho(r)dr,$$

where $V_{ion}(r)$ represents some local ionic potential function. The contribution from the non-local term is defined by

$$E_{ion(nonlocal)} = \sum_{i=1}^{k} \sum_{\ell} \left[\int \bar{\psi}_i(r) w_\ell(r) dr \right]^2,$$

where $w_\ell(r)$ denotes a pseudo-potential reference projection function.

The Hartree potential, defined by

$$V_H(r) = \int \frac{\rho(r')}{r - r'} dr',$$

is used to model the classical electrostatic average interaction between electrons. Its contribution to the total energy is defined by

$$E_H = \frac{1}{2} \int V_H(r)\rho(r)dr.$$

The exchange correlation function ϵ_{xc} is used to model the non-classical and quantum interaction between electrons. The potential energy induced by this

3

function is defined by

$$E_{XC} = \int \epsilon_{xc}(\rho(r))dr.$$

It is important to note that the minimization of the total energy (1) must be carried out under the orthonormality constraint

$$\int \bar{\psi}_i(r)\psi_j(r)dr = \delta_{i,j}.$$

In the following discussion, we will use A^T to denote the transpose of a matrix A, and A^* to denote the complex conjugate of A. A submatrix of A consisting of rows i through j and columns p through q will be denoted by the notation $A(i : j, p : q)$. If the submatrix contains all rows (columns) of A, it will be denoted by $A(:, p : q)$ ($A(i : j, :)$).

With an appropriate discretization scheme, the electron wave function $\psi_i(r)$ can be approximated by a vector $x_i \in \mathbb{C}^n$, where n is the spatial degree of freedom, i.e., the number of real space grid points. If we let $X = (x_1, x_2, ..., x_k)$, then the charge density associated with the k occupied states can be expressed by

$$\rho(X) = \text{diag}(XX^*),$$

where $\text{diag}(A)$ denotes a column vector consisting of diagonal entries of the matrix A.

Under the same discretization, the Laplacian operator ∇^2 can be approximated by a Hermitian matrix $L \in \mathbb{C}^{n \times n}$. The discretized local ionic potential can be represented by a diagonal matrix D_{ion}, and the discrete form of the Hartree potential can be represented by the product of a Hermitian matrix $S \in \mathbb{C}^{n \times n}$ with $\rho(X)$.

Using the notation established above, we can express the discrete form of various components of the total energy (1) by

$$E_{kinetic} = \frac{1}{2}\text{trace}(X^*LX), \tag{2}$$

$$E_{ion(local)} = \text{trace}(XD_{ion}X^*), \tag{3}$$

$$E_{ion(nonlocal)} = \sum_i \sum_\ell \left(x_i^* w_\ell\right)^2, \tag{4}$$

$$E_H = \frac{1}{2}\rho(X)^T S \, \rho(X), \tag{5}$$

$$E_{XC} = e^T \left(\epsilon_{xc}[\rho(X)] \right), \tag{6}$$

where e is a column vector of ones.

Once discretized, the minimization problem becomes

$$\begin{aligned} & \min \quad E_{total}(X) \\ & \text{s.t.} \quad X^*X = I_k, \end{aligned} \tag{7}$$

where I_k denotes a $k \times k$ identity matrix.

The Lagrangian associated with (7) is

$$\mathcal{L}(X) = E_{total}(X) - \text{trace}\left[\Lambda^T (X^*X - I_k) \right], \tag{8}$$

where Λ is a $k \times k$ matrix containing the Lagrange multipliers associated with the constraints specified by $X^*X = I_k$.

The solution to (7) must satisfy the first order necessary condition

$$\begin{aligned} & \nabla_X \mathcal{L}(X) = 0, \\ & \quad X^*X = I_k. \end{aligned} \tag{9}$$

Here, $\nabla_X \mathcal{L}$ represents an $n \times k$ matrix whose (i,j)-th entry is the partial derivative of $\mathcal{L}$ with respect to the (i,j)-th entry of X.

It is easy to verify that

$$\nabla_X E_{kinetic} = \frac{1}{2} LX, \tag{10}$$

$$\nabla_X E_{ion(local)} = D_{ion} X, \tag{11}$$

$$\nabla_X E_{ion(nonlocal)} = \sum_{\ell} (w_\ell w_\ell^*) X, \tag{12}$$

$$\nabla_X E_H = \text{Diag}(S\rho(X)) X, \tag{13}$$

$$\nabla_X E_{XC} = \text{Diag}(\mu_{xc}(\rho)) X, \tag{14}$$

where

$$\mu_{xc}(\omega) \equiv \frac{d\epsilon_{xc}(\omega)}{d\omega}$$

is the derivative of the exchange-correlation function. Here the notation $\text{Diag}(\rho)$ represents a diagonal matrix whose diagonal is determined by the vector ρ,

and we scaled (10) -(14) by 1/2 to be consistent with the convention used in the electronic structure community.

Substituting (10) - (14) into (9), we obtain the Kohn-Sham equation

$$\left[\frac{1}{2}L + D_{ion} + \sum_{\ell} w_{\ell}w_{\ell}^* + \text{Diag}(S\rho) + \text{Diag}(\mu_{xc}(\rho))\right]X = X\Lambda_k, \qquad (15)$$

$$X^*X = I_k. \qquad (16)$$

Because the vector ρ in (15) depends on X, the eigenvalue problem defined by (15) is nonlinear. Note that the solution to (7) is not unique. If X is a solution, then XQ is also a solution for some $Q \in \mathbb{C}^{k \times k}$ such that $Q^*Q = I_k$. That is, the solution to the constrained minimization problem or, equivalently, the nonlinear equations (15) is a k-dimensional invariant subspace in $\mathbb{C}^n$ rather than a specific matrix. In particular, Q can be chosen such that Λ_k is diagonal. In this case, X consists of k Kohn-Sham eigenvectors associated with the k smallest eigenvalues of (15).

3   The Self-Consistent Field Iteration

The most widely used method for computing the wave functions associated with the minimum total energy is the so-called *Self Consistent Field* (SCF) iteration, which corresponds to a fixed point iterative scheme applied to the Kohn-Sham equation (15). Given an initial guess of X, say $X^{(0)}$, one forms the discrete Hamiltonian,

$$H^{(1)} = \frac{1}{2}L + D_{ion} + \sum_{\ell} w_{\ell}w_{\ell}^* + \text{Diag}(S\rho(X^{(0)})) + \text{Diag}(\mu_{xc}(\rho(X^{(0)}))), (17)$$

and computes eigenvectors $X^{(1)}$ associated with the k smallest eigenvalues of $H^{(0)}$. These eigenvectors defines a new Hamiltonian $H^{(2)}$. In the basic version of the SCF iteration, the difference between $\rho(X^{(i-1)})$ and $\rho(X^{(i)})$ is examined at each step to determine whether the iteration should be terminated. If the change in the charge density remains large, the eigenvectors associated with the k smallest eigenvalues of $H^{(i)}$ are computed, and this process continues until the $\|\rho(X^{(i-1)}) - \rho(X^{(i)})\|$ becomes negligibly small. In this case, $X^{(i)}$ consists of a set of wave functions that are self-consistent with respect to the KS equation (15).

For completeness, we outline the major steps of the basic version of a SCF calculation in Figure 1. Depending on the discretization scheme used, it may

```
  SCF Iteration

  Input:    The matrices $L$, $D_{ion}$, $S$, the vectors $w_\ell$, $\ell = 1, 2, ....$ The
            derivative of the exchange-correlation function $\mu_{xc}(x)$; an
            initial guess $X^{(0)}$ for the optimal wave function $X \in \mathbb{C}^{n \times m}$;

  Output:   $X \in \mathbb{C}^{n \times m}$ such that $X^* X = I_m$ and $E_{tot}(X)$ is minimized,
            where $E_{tot}(X)$ is defined by (10) - (14).



  1.        for $i = 1, 2, ...$ until convergence
  2.          Form $H^{(i)} = H(X^{(i-1)})$;
  3.          Compute $X^{(i)}$ such that $H^{(i)} X^{(i)} = X^{(i)} \Lambda^{(i)}$, and $\Lambda^{(i)}$
              contains the $k$ smallest eigenvalues of $H^{(i)}$;
  4.        end for
```

Fig. 1. The SCF iteration

not be necessary or possible to form the Hamiltonian $H^{(i)}$ explicitly in the
SCF calculation. This is particularly true when the continuous problem is
discretized by a spectral method using a plane wave basis. In that case, $H^{(i)}$
only exists in the form of matrix vector multiplication procedure.

For large-scale problems or problems in which $H^{(i)}$ cannot be formed explic-
itly, it is usually not feasible to solve the linear eigenvalue problem $H^{(i)} X^{(i)} =
X^{(i)} \Lambda^{(i)}$ by using a QR (Francis 1961; Francis 1962) type of eigensolver. It-
erative methods such as the Lanczos (Lanczos 1950), preconditioned conju-
gate gradient (Hestenes & Karush 1951) or a Jacobi-Davidson type of method
(Sleijpen & Van der Vorst 1996; Davidson 1975; Olsen et al. 1990) are more
appropriate in this setting.

The major computational cost of the SCF scheme is in solving a linear eigen-
value problem at Step 3 in each iteration. Because the condition number of
$H^{(i)}$ is typically large ($> 10^6$), and the eigenvalues of interest are typically
clustered at the low end of the spectrum, solving such an eigenvalue problem
can be extremely challenging. In the materials science community, the com-
mon practice for computing the desired eigenvalues of $H^{(i)}$ is to apply a pre-
conditioned conjugate gradient algorithm to minimize the Rayleigh Quotient
$u^* H^{(i)} u / u^* u$. Eigenvalues are often computed one at a time, and standard
deflation techniques are employed to ensure the computed eigenvectors are
mutually orthogonal.

Although an SCF iteration is often viewed as a fixed-point iteration, one cannot write down an explicit mapping from $X^{(i-1)}$ to $X^{(i)}$ even if the linear eigenvalue problem

$$H^{(i)}X^{(i)} = X^{(i)}\Lambda^{(i)}, \tag{18}$$

can be solved exactly. Therefore, it is somewhat difficult to analyze the convergence of the SCF iteration using the standard theory for a fixed point iteration without making certain assumptions.

Although the basic SCF iteration works on some problems, it tends to fail for most large systems. In practice, a modified scheme in which a mixture of $H^{(i)}$ and $H^{(i-1)}$ is used to replace $H^{(i)}$ in Step 2 of the algorithm is much more effective. This scheme is called *charge* or *potential mixing* (Kresse & Furthmüller 1996b; Kresse & Furthmüller 1996a). Furthermore, charge mixing is often combined with the use of the direct inversion of iterative subspace (DIIS) algorithm proposed in (Pulay 1980; Pulay 1982) to accelerate the convergence of SCF. However, there is no theoretical guarantee that these techniques will always work. In fact, there are cases in which these techniques fail also.

4   A Constrained Optimization Algorithm for Total Energy Minimization

The algorithm we present in this section aims at minimizing the KS total energy functional directly. This general approach has been discussed in a number of papers (Gillan 1989; Teter 1989; Payne et al. 1992; Kresse & Furthmüller 1996a; Voorhis & Head-Gordon 2002). In (Teter 1989; Payne et al. 1992), a conjugate gradient (CG) type of algorithm is used to minimize the total energy. The minimization is carried out "band-by-band", i.e., the total energy is minimized with respect to one wave function at a time. For the j-th band (wavefunction), the search direction $p_j^{(i)}$ is generated from a linear combination of the wavefunction $x_j^{(i)} = X^{(i)}e_j$ and the residual

$$r_j = H^{(i)}x_j^{(i)} - x_j^{(i)}\lambda_j,$$

where λ_j is the j-th eigenvalue of the projected Hamiltonian $X^{(i)^*}H^{(i)}X^{(i)}$. Note that r_j is simply the j-th column of the gradient matrix $\nabla_X \mathcal{L}(X^{(i)})$. Similar to a standard CG algorithm, the linear combination of $x_j^{(i)}$ and r_j is chosen so that $p_j^{(i)}$ is $H^{(i)}$-conjugate to the previous search direction $p_j^{(i-1)}$. The new wavefunction $x_j^{(i+1)}$ is then computed by minimizing the KS total energy in the subspace spanned by $x_j^{(i)}$ and $p_j^{(i)}$. To simplify this minimization problem,

8

$p_j^{(i)}$ is first orthogonalized against $x_j^{(i)}$ and normalized so that $\|p_j^{(i)}\| = 1$. The new wavefunction is parameterized by

$$x_j^{(i+1)} = x_j^{(i)} \cos\theta + p_j^{(i)} \sin\theta,$$

where the optimal θ is obtained by a standard line search procedure. Instead of using the KS functional to perform the line search, Teter et al. (Teter 1989) proposed using a surrogate function that is cheaper to evaluate. However, this approach was shown in (Kresse & Furthmüller 1996a) to be less efficient than the SCF iteration. We believe this is primarily due to the "band-by-band" nature of the algorithm.

A constrained Quasi-Newton algorithm is used in (Voorhis & Head-Gordon 2002) to minimize the total energy with respect to all wave functions (associated with the occupied states) simultaneously. The algorithm first computes the search direction via a limited-memory BFGS (Liu & Nocedal 1989) scheme, the search direction is then modified through a parallel transport technique (Edelman et al. 1998) to ensure that the orthonormality constraint $X^*X = I_k$ is satisfied in the line search procedure. It is reported in (Voorhis & Head-Gordon 2002) that this constrained Quasi-Newton approach is generally slower than the SCF iteration combined with DIIS extrapolation. It was also found that the optimization procedure can sometimes be trapped at a local minima. It was conjectured that these problems are associated with the poor approximation to the Hessian of the total energy function due to the large number of degrees of freedom. However, there are cases in which the constrained Quasi-Newton algorithm converges to the optimal solution while the SCF iteration fails to converge.

The direct minimization algorithm we present here also seeks the optimal wave functions associated with all occupied states simultaneously as in (Voorhis & Head-Gordon 2002). However, we choose the search direction from a subspace that consists of the existing wave functions $X^{(i)}$, the gradient of the Lagrangian (8) and the search direction computed in the previous iteration. A special line search strategy is developed to minimize the total energy within the search space while maintaining the orthonormality constrained required for $X^{(i+1)}$. This strategy requires us to solve a projected nonlinear eigenvalue problem as we will show below.

Let $R^{(i)}$ be the preconditioned gradient of the Lagrangian (8) with respect to X evaluated at $X^{(i)}$, and let $P^{(i-1)}$ be the search direction obtained in the i-1st iteration. In our algorithm, the wave function update is performed within the $3k$-dimensional subspace spanned by $X^{(i)}$, $R^{(i)}$ and $P^{(i-1)}$. If we let

$$Y = (X^{(i)}, \ R^{(i)}, \ P^{(i-1)}),$$

we can then express the new approximation, $X^{(i+1)}$, by

$$X^{(i+1)} = YG, \tag{19}$$

where $G \in \mathbb{C}^{3k \times k}$ is chosen to minimize $\hat{E}(G) \equiv E_{total}(YG)$, i.e. we must solve

$$\begin{aligned} &\min_G E_{total}(YG) \\ &\text{s.t. } G^*Y^TYG = I_k. \end{aligned} \tag{20}$$

The first order necessary condition of (20) can be derived by examining the gradient of $\hat{E}(G)$ with respect to G.

It is easy to verify that

$$\begin{aligned} \nabla_G \hat{E}_L(G) &= \frac{1}{2}\nabla_G\Big[\text{trace}(G^*Y^*LYG)\Big] \\ &= (Y^*LY)G, \end{aligned} \tag{21}$$

$$\begin{aligned} \nabla_G \hat{E}_V(G) &= \nabla_G\Big[\text{trace}(G^*Y^*D_{ion}YG)\Big] \\ &= Y^*D_{ion}YG, \end{aligned} \tag{22}$$

$$\begin{aligned} \nabla_G \hat{E}_W(G) &= \nabla_G\Big[\sum_i \sum_\ell (w_\ell^* Y g_i)^2\Big] \\ &= \sum_\ell (Y^*w_\ell)(Y^*w_\ell)^*G, \end{aligned} \tag{23}$$

$$\begin{aligned} \nabla_G \hat{E}_R(G) &= \frac{1}{2}\nabla_G\Big[\rho(YG)^T S\rho(YG)\Big] \\ &= Y^*\text{Diag}\Big[S\rho(YG)\Big]YG, \end{aligned} \tag{24}$$

$$\begin{aligned} \nabla_G \hat{E}_X(G) &= \nabla_G\Big[\epsilon_{xc}(\rho(G))\Big] \\ &= Y^*\text{Diag}\Big[\mu_{xc}(\rho(YG))\Big]YG. \end{aligned} \tag{25}$$

Again, (21) - (25) have been scaled by $1/2$ to be consistent with the convention used in the electronic structure community.

If we define

$$\hat{H}(G) = Y^*\Big[\frac{1}{2}L + D_{ion} + \sum_\ell w_\ell w_\ell^* + \text{Diag}\Big(S\rho(YG)\Big) + \text{Diag}\Big(\mu_{xc}(\rho(YG))\Big)\Big]Y, \tag{26}$$

then, solving (20) is equivalent to solving

$$\hat{H}(G)G = BG\Omega_k, \tag{27}$$
$$G^*BG = I_k, \tag{28}$$

where $B = Y^*Y$ and the $k \times k$ diagonal matrix Ω_k contains the k smallest eigenvalues of (27).

Note that the projected nonlinear eigenvalue problem defined by (27) and (28) is much smaller than the nonlinear eigenvalue solved in an SCF iteration. The reduction in size provides us with more flexibility in terms of the algorithms we can choose to solve the nonlinear eigenvalue problem. For example, if we apply an SCF iteration to compute the desired eigenpairs of (27) and (28), the linear eigenvalue problem that emerges from each SCF iteration can be solved by using the LAPACK (Anderson et al. 1992) implementation of the QR algorithm (Francis 1961; Francis 1962). We may even apply a standard nonlinear constrained minimization algorithm such as a sequential quadratic programming technique (SQP) (Boggs & Tolle 1995) to (20) directly. Furthermore, it should be noted that it is not necessary to solve equations (27) - (28) to full accuracy in the early stage of the direct minimization process because all we need is a G that yields sufficient decrease in the objective function within the subspace spanned by columns of Y.

Once G is computed, we can update the wave function following (19). In addition, we can compute the search direction associated with this update

$$P^{(i)} \equiv X^{(i+1)} - X^{(i)}G(1:k,:) = Y(:,k+1:3k)G(k+1:3k,:).$$

Because the solution to (27)-(28) ensures columns of $X^{(i+1)}$ are orthonormal, there is no need to explicitly orthogonalize $P^{(i)}$ against $X^{(i)}$ in our algorithm.

A complete description of the constrained minimization algorithm is shown in Figure 2. We should point out that solving the projected optimization problem in Step 6 of the algorithm may require us to evaluate the projected Hamiltonian (26) repeatedly as we search for the best G. However, since the first three terms of $\hat{H}$ do not depend on G, they can be computed and stored in advance. Only the last two terms of (26) need to be updated. These updates requires the charge density, the Hartree and the exchange-correlation potentials to be recomputed.

5 Numerical Example

In this section, we compare the performance of the direct constrained minimization (DCM) algorithm presented in the previous section with that of the

Algorithm: A Constrained Minimization Algorithm for Total Energy Minimization

Input: An initial set of wave function $X^{(0)} \in \mathbb{C}^{n \times k}$; the matrices L, D_{ion}, S; the vectors w_ℓ, $\ell = 1, 2, ...$ The derivative of the exchange-correlation function $\mu_{xc}(x)$; a preconditioner K;

Output: $X \in \mathbb{C}^{n \times k}$ such that the KS total energy functional $E_{total}(X)$ is minimized and $X^*X = I_k$.

1. Orthonormalize $X^{(0)}$ such that $X^{(0)^*}X^{(0)} = I_k$;

2. for $i = 0, 1, 2, ...$ until convergence

3. Compute $\theta_j = e_j^T X^{(i)^*} H^{(i)} X^{(i)} e_j$, $j = 1, 2, ..., k$;

4. Compute $R = K^{-1}\left[H^{(i)}X^{(i)} - X^{(i)}D \right]$,

   where $D = \text{Diag}(\theta_1, \theta_2, ..., \theta_m)$;

5. if $(i > 1)$ then

   $Y \leftarrow (X^{(i)}, R, P^{(i-1)})$

   else

   $Y \leftarrow (X^{(i)}, R)$;

   endif

6. $B \leftarrow Y^*Y$;

7. Find $G \in \mathbb{C}^{2k \times 2k}$ or $\mathbb{C}^{3k \times 3k}$ that minimizes $E_{total}(YG)$

   subject to the constraint $G^*BG = I$;

8. Set $X^{(i+1)} = YG$;

9. if $(i > 1)$ then

   $P^{(i)} \leftarrow Y(:, k+1:3k)G(k+1:3k,:)$;

   else

   $P^{(i)} \leftarrow Y(:, k+1:2k)G(k+1:2k,:)$;

   endif

10. end for

Fig. 2. A Direct Constrained Minimization Algorithm for Total Energy Minimization

SCF iteration implemented in the software package PEtot (Wang) through a numerical example. In PEtot, $\psi_i(r)$ is discretized by a spectral method using plane waves as the basis. These basis functions are eigenfunction of the Laplacian operator (L) associated with the kinetic energy of the atomistic system. Thus, PEtot stores only the Fourier coefficients of each wave function $x_j = Xe_j$ instead of x_j itself so that $y \leftarrow Lx_j$ can be carried out in $\mathcal{O}(n)$ floating point operations (flops) in the frequency domain. However, because the potential terms of the Hamiltonian (with the exception of the non-local ionic potential) are diagonal in the spatial domain, PEtot converts the Fourier space representation of x_j into the real space representation before operations involving these potential terms are performed. The complexity of this conversion is $\mathcal{O}(n \log n)$ when it is carried out by a Fast Fourier Transform (FFT).

We measure the convergence of both algorithms by examining the relative reduction of the total energy computed in each outer iteration. The relative reduction is evaluated by

$$\Delta E_i = E_{total}(X^{(i)}) - E_{min},$$

where E_{min} is a lower bound of the total energy.

The computational cost of both the SCF and DCM is estimated by the number of matrix vector multiplications $y \leftarrow H^{(i)}x$ performed in these two algorithms. To illustrate that this is a reliable measure when $k \ll n$, we also provide the timing measurements. Our computation is performed on a single node of the IBM SP maintained at NERSC. Each IBM SP node contains 16 Power3 CPUs and 16 GB memory. Each Power3 CPU runs at a 375Mhz clock speed, and has 2 MB L2 cache. Both SCF and DCM are parallelized using MPI. We used the IBM math library ESSL for dense matrix and FFT calcluations.

In a PEtot SCF iteration, the linear eigenvalue problem (18) is solved by applying a preconditioned conjugate gradient (PCG) algorithm to minimize the Rayleigh quotient $x^*H^{(i)}x/x^*x$. Explicit deflation is put in place to ensure the convergence of the smallest k eigenpairs. Each PCG iteration requires a single matrix vector (MATVEC) multiplication followed by a preconditioning operation. When n is sufficiently large, the complexity of each MATVEC is dominated by the cost of the FFT calculation used to convert the Fourier space representation of x_j to the real space representation. The Laplacian operator L is used as the preconditioner. Because it is diagonal in the frequency space, the cost of preconditioning is relatively small compared to a MATVEC. If m PCG iterations are taken on average to compute an approximate eigenpair of (18), then the total number of MATVECs used per SCF iteration is $m \times k$.

In the direct constrained minimization (DCM) algorithm, k MATVECs are performed in each outer iteration to compute the gradient. The inner iteration

13

does not require a full MATVEC in the form of $y \leftarrow H^{(i)}x$. However, the update of the projected Hartree potential at each inner iteration requires us to compute $S\rho(YG)$. Because S is the inverse of L, this calculation is typically carried out by a fast Poisson solver. The complexity of this computation is approximately $\mathcal{O}(n \log n)$, which is equivalent to a single MATVEC used in the SCF iteration asymptotically. Thus, if p inner SCF iterations are taken in the DCM algorithm, the total number of MATVECs used per DCM iteration is $k + p$.

We should point out that both the SCF and the DCM algorithms perform an additional $\mathcal{O}(n \cdot k^2)$ operations in each inner iteration. In SCF, these operations were used to maintain the orthogonality of the Fourier representation of x_j. In DCM, these operations were used to update the charge density $\rho(YG)$ and to compute the projected potential term

$$Y^* \left[\mathrm{Diag}\Big(S\rho(YG) + \mu_{xc}(\rho(YG)) \Big) \right] Y.$$

The cost of these operations is relatively small compared to that associated with FFT when $n \gg k$. However, when k becomes a significant fraction of n, the cost of these calculations cannot be ignored.

We applied both algorithms to a simple SiH4 test problem. The problem is discretized on a $32 \times 32 \times 32$ real space grid. The number of plane wave basis functions used in the Fourier representation is 2103. The number of occupied states for this molecule is $k = 4$.

In the SCF calculation, we set the convergence tolerance of each PCG run to $\tau = 10^{-12}$ and the maximum number of PCG iterations allowed to 10. That is, we terminate the PCG iteration when

$$\|H^{(i)}x_j^{(i)} - \lambda_j^{(i)}x_j^{(i)}\| \leq 10^{-12},$$

or when the number of PCG iterations taken reaches 10. In our experiment, the PCG convergence tolerance was never reached before the maximum number of iterations were taken. Thus each outer SCF iteration consumed 4×10 MATVECS. Both Pulay-Kerker and Pulay-ThomasFermi charge mixing scheme were used in the outer SCF iteration to accelerate the convergence.

In DCM, the projected minimization problem was solved by applying a simple SCF iteration (without charge mixing) to (27). We set the number of inner SCF iterations to 3.

Figure 3 shows that both SCF and DCM reached the same total energy level after 15 outer iterations. Clearly, DCM consumed a much smaller number of

MATVECs. Furthermore, we observed that the reduction in total energy is monotonic in DCM. In the PEtot version of SCF, the KS Hamiltonian was not updated until the end of 4-th outer iteration. This was required to carry out the charge mixing procedure for accelerating the convergence of SCF. As a result, the total energy objective does not show significant improvement between the 2nd and the 4-th outer iteration.
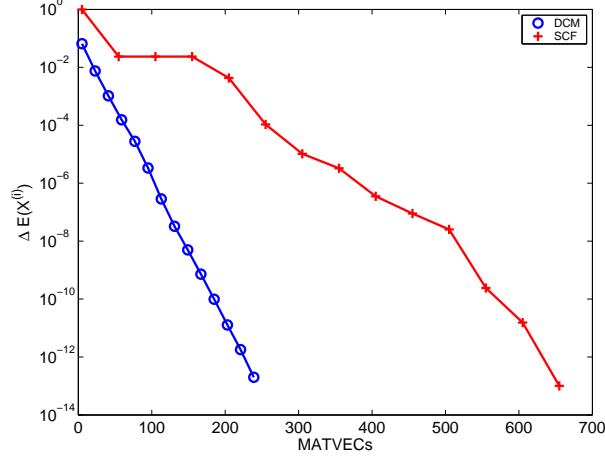

Fig. 3. Comparing the convergence of SCF and DCM in terms of the number of MATVECs performed.

Because the number of occupied states ($k = 4$) in SiH4 is relatively small, both SCF and DCM are dominated by the FFT computation required in each MATVEC. Hence, the MATVEC count shown in Figure 3 gives a good measure of how the two methods compare in terms of computational speed. In Figure 4 we plot the reduction of the total energy with respect to the wall clock time used in both SCF and DCM. The figure shows that DCM is almost 4 times faster than SCF in terms of wall clock time.
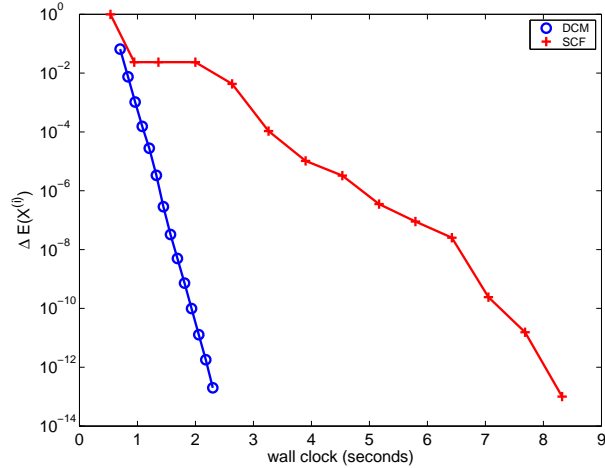

Fig. 4. Comparing the performance of DCM with SCF in terms of wall clock time.

6   Concluding Remarks

A direct constrained minimization (DCM) algorithm for computing the ground state total energy of a large scale atomistic system is presented in this paper. The algorithm constructs a new search direction from the subspace $\mathcal{S}$ spanned by the current approximation to the optimal wavefunction, the gradient of the total energy and the previous search direction. The optimal search direction and step length are computed by minimizing the total energy functional within $\mathcal{S}$ subject to an orthonormality constraint. Solving this smaller optimization problem is equivalent to solving a small nonlinear eigenvalue problem. In our computational scheme, all wave functions associated with the occupied states are updated simultaneously instead of "band-by-band".

We compared the convergence of the DCM algorithm with that of the SCF iteration implemented in PEtot through a numerical example. We demonstrated that DCM can be more efficient than SCF. Because DCM modifies the wave functions associated with all occupied states simultaneously (a block algorithm), we can make better use of the cache and memory in each MATVEC thereby further improving the performance of the computation.

It is observed that neither DCM nor SCF requires each inner iteration to converge to full accuracy in order to reach the optimal solution. In our numerical experiment, we set the number of inner iterations in both methods to a fixed number. However, this is clearly not the best strategy. More research is required to develop a more effective stopping criterion for the inner iteration in both methods.

In our numerical experiment, the projected optimization problem (20) is solved by applying a basic SCF iteration to (27). This simple scheme worked well for the SiH4 system. However, it is conceivable that the inner SCF iteration may fail to converge or fail to produce a sufficient reduction in the objective function for other systems. It may be necessary to use the DIIS or other charge mixing extrapolation scheme to improve the convergence of the inner iteration. It may also be possible to replace the SCF iteration with a Newton type of algorithm because of the small dimension of the projected problem.

References

[Anderson et al. 1992] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide.* SIAM, Philadelphia, PA., second edition, 1992.

[Bendt & Zunger 1982] P. Bendt and A. Zunger. New approach for solving the

density-functional self-consistent field problem. *Physical Rev. B*, 26:3114–3137, 1982.

[Boggs & Tolle 1995] P. T. Boggs and J. W. Tolle. Sequential quadratic programming. *Acta Numerica*, pages 1–52, 1995.

[Davidson 1975] E. R. Davidson. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real symmetric matrices. *Journal of Computational Physics*, 17:87–94, 1975.

[Edelman et al. 1998] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20(2):303–353, 1998.

[Francis 1961] J. G. F. Francis. The QR transformation a unitary analogue to the LR transformation - part 1. *The Computer Journal*, 4:265–271, October 1961.

[Francis 1962] J. G. F. Francis. The QR transformation - part 2. *The Computer Journal*, 4:332–345, January 1962.

[Gillan 1989] M. J. Gillan. Calculation of the vacancy formation in aluminum. *J. Phys. Condens. Matter*, 1:689–711, 1989.

[Goedecker & Scuseria 2003] S. Goedecker and G. Scuseria. Linear scaling electronic structure methods in chemistry and physics. *IEEE Comp. Sci. Eng.*, 5:14–21, 2003.

[Hestenes & Karush 1951] M. R. Hestenes and W. Karush. A method of gradients for the calculation of the characteristic roots and vectors of a real symmetric matrix. *Journal of Research of the National Bureau of Standards*, 47:45–61, 1951.

[Kresse & Furthmüller 1996a] G. Kresse and J. Furthmüller. Efficiency of ab initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Computational Materials Science*, 6:15–50, 1996.

[Kresse & Furthmüller 1996b] G. Kresse and J. Furtthmüller. Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Physical Review B*, 54(16):11169 – 11185, 1996.

[Lanczos 1950] C. Lanzcos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of Research of the National Bureau of Standards*, 45(4):255–282, October 1950. Research Paper 2133.

[Liu & Nocedal 1989] D. C. Liu and J. Nocedal. On the limited memory method for large scale optimization. *Math. Programming, B*, 45(3):503–528, 1989.

[Olsen et al. 1990] J. Olsen, P. Jorgensen, and J. Simons. Passing the one-billion limit in full configuration-interaction (fci) calculations. *Chemical Physics Letters*, 169:463–472, 1990.

[Payne et al. 1992] M. C. Payne, M. P. Teter, D. C. Allen, T. A. Arias, and J. D. Joannopoulos. Iterative minimization techniques for *ab initio* total energy calculation: molecular dynamics and conjugate gradients. *Reviews of Modern Physics*, 64(4):1045–1097, 1992.

[Pfrommer et al. 1999] B. G. Pfrommer, J. Demmel, and H. Simon. Uncon-

strained Energy Functionals for Electronic Structure Calculations. *J. Comp. Phys.*, 150:287–298, 1999.

[Pulay 1980] P. Pulay. Convergence acceleration of iterative sequences: The case of SCF iteration. *Chemical Physics Letters*, 73(2):393–398, 1980.

[Pulay 1982] P. Pulay. Improved SCF convergence acceleration. *Journal of Computational Chemistry*, 3(4):556–560, 1982.

[Sleijpen & Van der Vorst 1996] G. L. G. Sleijpen and H.A. Van der Vorst. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM Journal on Matrix Analysis and Applications*, 17(2):401–425, April 1996.

[Teter 1989] M. P. Teter, M. C. Payne, and D. C. Allan. Solution of Schrödinger's equation for large systems. *Physical Review B*, 40(18):12255–12263, 1989.

[Voorhis & Head-Gordon 2002] T. Van Voorhis and M. Head-Gordon. A geometric approach to direct minimization. *Molecular Physics*, 100(11):1713–1721, 2002.

[Wang] L. Wang. Parallel Total Energy Calculation Software. http://hpcrd.lbl.gov/~linwang/PEtot/PEtot.html.